



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 7 : H04Q 11/04, H04L 12/56		A1	(11) International Publication Number: WO 00/47011 (43) International Publication Date: 10 August 2000 (10.08.00)
<p>(21) International Application Number: PCT/US00/02604</p> <p>(22) International Filing Date: 31 January 2000 (31.01.00)</p> <p>(30) Priority Data: 09/241,597 2 February 1999 (02.02.99) US</p> <p>(71) Applicant: CISCO TECHNOLOGY, INC. [US/US]; 170 West Tasman Drive, San Jose, CA 95134 (US).</p> <p>(72) Inventors: FOURIE, Henry; 16416 Shadyview Lane, Los Gatos, CA 95032 (US). KARIA, Snehal, G.; 1026 Dolphin Common, Freemont, CA 94536 (US).</p> <p>(74) Agents: VINCENT, Lester, J. et al.; Blakely, Sokoloff, Taylor & Zafman, 12400 Wilshire Boulevard, 7th Floor, Los Angeles, CA 90025-1026 (US).</p>		<p>(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).</p> <p>Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i></p>	
<p>(54) Title: SWITCHED VIRTUAL CIRCUIT CONTROLLER FOR A SETUP CONGESTION MANAGEMENT STRATEGY</p> <p>(57) Abstract</p> <p>A method that specifies a congestion management strategy for a SVC (Switched Virtual Circuit) controller in a connection-oriented network. The rate of setup messages arriving at a given interface handled by the controller and the aggregate number of calls being established are compared against a set of thresholds to determine whether or not to allow new call setup messages from being processed. The congestion strategy regulates the consumption of controller resources such as processor load and memory utilization.</p>			
<pre> graph TD 610[BEGIN INTERVAL] --> 615[MONITOR SETUP ARRIVAL RATE] 615 --> 625{SETUP ARRIVAL RATE > UPPER SETUP THRESHOLD?} 625 -- YES --> 650{AGGREGATE CALLS IN PROGRESS > UPPER CALLS THRESHOLD?} 650 -- YES --> 655[ONSET OF CONGESTION] 655 --> 660[DISCARD SETUP MESSAGES] 660 --> 645[SETUP MESSAGES ALLOWED] 645 --> 640[NO CONGESTION/ABATEMENT OF CONGESTION] 640 --> 610 625 -- NO --> 635{AGGREGATE CALLS IN PROGRESS < LOWER CALLS THRESHOLD?} 635 -- YES --> 640 635 -- NO --> 630{SETUP ARRIVAL RATE < LOWER SETUP THRESHOLD?} 630 -- NO --> 640 630 -- YES --> 610 </pre>			

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

-1-

SWITCHED VIRTUAL CIRCUIT CONTROLLER FOR A SETUP CONGESTION MANAGEMENT STRATEGY

1. Field of the Invention

The invention relates generally to communications and networking. More specifically, the invention relates to the usage of resources in networking devices.

2. Background of the Invention

In connection-oriented networking schema such as ATM (Asynchronous Transfer Mode), connections or "calls" must be established between one information device such as a computer system or router and another. This call or connection is sometimes referred to as a "virtual circuit" (VC) particularly where a specified data pipe is artificially, through software, segmented into separate data-pathways, each pathway servicing a particular VC. Often a switch acts as an intermediary to direct one or more of these VCs through a particular network node, and thus these calls are collectively referred to as SVCs (Switched Virtual Circuits).

Figure 1 shows an exemplary wide-area networking system serviced by ATM. A wide-area network (WAN) link 120 interconnects a first network 100 with a second network 110. Each network has a plurality of nodes which may each contain switching devices that regulate data traffic to one or more user terminals.

-2-

Network 100 is shown having nodes 102, 103, 104, 106 and 108, while network 110 is shown having nodes 112, 114, 116 and 118. A first user terminal 105 is connected to node 102 of network 100 while a second user terminal 115 is connected to node 118 of network 110. In order for user terminal 105 and user terminal 115 to communicate with one another, a call must first be established between them. This call may be switched through a plurality of nodes. One possible route for sending data from user terminal 105 to user terminal 115 is for data to go from node 102 to 106 to 108 and then across the WAN link to node 112 and node 118 finally reaching user terminal 115.

Similar to PSTN (Public Switched Telephone Network) communications such as telephone calls, the period of SVC call operation for a given call can be split into three distinct phases—establishment (setup), active (data transfer) and disconnect (hang-up). Once a call is established, say between user terminal 105 and user terminal 115 across a specified path, a virtual circuit will have been created and the call can proceed into the active phase where data is transferred. Once the data transfer is complete, the call can be disconnected which will release the virtual circuit.

Each node has a controller device (SVC controller) and switch which facilitates the calls through its node. The SVC controller has processing, memory and other resources to interpret, forward and

-3-

process messages and initiate other messages as appropriate, while the switch ordinarily handles the physical routing of messages among nodes and user terminals.

In a wide-area network, as that shown in **Figure 1**, call setup may have to proceed through an arbitrarily large and unknown number of nodes before it can be completed. **Figure 2** shows a call setup session between a first user terminal and a second user terminal that proceeds through a series of N nodes. Each of the nodes $1, 2, \dots, N$ have a controller device that performs, among other functions, the data link layer (Layer 2) and transport layer (Layer 3) processing and forwarding of messages. Among such messages are call setup messages. For instance, in **Figure 2**, when establishing a call between user terminal 1 and user terminal 2, user terminal 1 will first initiate a setup message informing the node 1 controller that a connection wishes to be established with user terminal 2. The node 1 controller sends back an acknowledgement message informing user terminal 1 that the call establishment is still in progress (call proceeding). The node 1 controller will then forward the call setup initiator message to next down-the-line node, node 2. Upon receiving the setup message, the node 2 controller responds with a call proceeding message to the node 1 controller. The node 2 controller forwards the call setup message to another node's controller, which

-4-

responds with a call proceeding message to node 2, and on and on, until reaching the second user terminal 2 which is directly accessible from node N. User terminal 2 sends a call proceeding message back to the node N controller, and after a time, when the call is connected through, the user terminal 2 sends to the node N controller a call connected message which is forwarded back through the nodes numbered one to N until reaching user terminal 1. At this point, the call is considered "established" and the data transfer phase commences.

The above exemplifies the number of messages needed to setup a single call over N nodes. Each node's controller would have to process and forward at least three messages to facilitate call establishment. Assuming each node is capable of supporting/switching multiple user terminals or other devices, many unique point-point calls are possible. A particular node controller may thus have to process many call setup messages. Each SVC controller has finite resources to provide these and other services, namely, message buffers, call records, CPU processing capacity, I/O bandwidth and so on which are often taxed. Thus, a high setup message arrival rate can potentially consume all of a controller's processing and memory resources. This can result in the controller delivering degraded service to all user terminals. A high setup message arrival rate due to one terminal

-5-

can thus act to degrade the service provided to all other user terminals.

SUMMARY OF THE INVENTION

What is disclosed is a method of managing congestion in a network controller having a plurality of interfaces comprising monitoring the utilization and potential utilization of resources in a due to message traffic received by the network controller, and instituting controls to limit further message traffic where the measurements exceed upper thresholds.

BRIEF DESCRIPTION OF THE DRAWINGS

The objects, features and advantages of the method and apparatus for the present invention will be apparent from the following description, wherein like references indicate similar elements, in which:

Figure 1 illustrates an exemplary networking topology.

Figure 2 illustrates the messages and their forwarding in a call establishment.

Figure 3 is a block diagram of call congestion monitoring in a SVC controller.

Figure 4 illustrates a methodology for tracking the aggregate number of calls in the establishment phase.

-6-

Figure 5 is a control and message flow diagram of call congestion monitoring and management.

Figure 6 is a flowchart of call congestion management.

DETAILED DESCRIPTION

What is described is a congestion management strategy for a network controller device such as an SVC (Switched Virtual Circuit) controller that interprets and processes messages related to connection-oriented network calls transported over interfaces connected to the controller. The strategy minimizes call setup traffic with the intended advantage of maintaining a certain level of fairness and quality of service to calls already connected when congestion situations arise.

The strategy involves first determining whether or not the system is congested, primarily, by deriving two metrics, which will be described in more detail below, and comparing these metrics to a set of thresholds. A control is instituted that, upon a determination of congestion disallows future setup messages. The intended advantages of doing so are that calls that are already established are delivered more reliable service and crashes due to resource overload are minimized.

Figure 3 is a block diagram of call congestion monitoring in a SVC controller.

-7-

Call setup congestion management can be decomposed into two operations: a measurement operation and a control operation. **Figure 3** illustrates a network node that contains a switch device 320 and a SVC controller 300 connected together. The SVC controller has at least two functions, layer 3 and layer 2, which refer to the network and data link layers in the well-known OSI (Open Systems Interconnect) networking model. Switch 320 interconnects the node that it resides in with other nodes in the network, while the SVC controller encodes/decodes and processes all the messages that arrive at the node, and additionally, is responsible for establishing signaling channels with user terminals connected to the switch and the controller itself. Switch 320 also connects user terminals that are connected to the node. Each of the physical connections from the switch to the SVC controller 300, whether to a user terminal or to another node, is referred to hereinafter as an "interface" upon which congestion control is applied. A single interface may host more than "virtual circuit", based upon bandwidth and configuration of the network. When messages of any type are received at the node at a particular interface, they first pass through switch 320 and then are forwarded to the layer 2 function of controller 300. The layer 2 function of controller 300 is also referred to as SSCOP (Service-Specific Connection-Oriented Protocol, set forth in the ITU (International Telecommunication Union) specifications Q.2110 and

-8-

Q.2130) which provides a reliable data link transport function to layer 3. The layer 3 function performs both signaling and call processing. Not shown is an intermediary buffer between layer 2 and layer 3 which is incorporated for sequencing and timing in some systems.

In the art currently, the SSCOP has a windowing mechanism that can be used to control the flow of signaling messages. This flow control mechanism applies to all signaling messages regardless of content and thus regardless of their potential usage or effect on controller resources. Under the windowing mechanism, once signaling messages become too burdensome, all further signaling messages are blocked. Certain signaling messages, such as disconnect phase messages, however, would be useful to process at layer 3 since they may have a net effect of freeing system resources. To prevent this arbitrary windowing, an enhanced congestion management strategy has been developed that behaves in accordance with the type of message received. The congestion strategy regulates the consumption of controller resources such as processor load and memory utilization but does so interface by interface so that a single user (on that interface) does not act to degrade quality of service provided to all other users, particularly those whose calls are already being established or are already active. By doing so, a level of fairness is delivered to all users.

-9-

Messages that free system resources even though initially requiring processing activity by the CPU are the "Release", "Release Complete", "Drop Party", "Drop Party Ack" and "Drop Party Reject" messages. Although the exact nomenclature and number of signaling messages will vary from system to system, these messages, whatever form they may take can be prioritized as follows:

1. Messages that complete release phase activity thus freeing system resources.
2. Messages that use system resources but initiate the call release phase, thus leading to an eventual freeing of resources.
3. Messages that complete call establishment, since call setup initiation is already complete, in fairness, the call should be allowed to be connected through.
4. Messages that initiate call establishment consuming resources without guarantee that a call will be completed i.e., they are risky in comparison to the benefit provided.

Using the above classification, messages of the first type have the highest priority, type 2 are second in priority, type 3 are next and finally messages of the fourth type, the lowest in priority. The lower the priority of the message, the more effective it is to flow-control them.

-10-

The two primary metrics for determining when and what controls should be employed even given the above classification are 1) a Setup message arrival rate for each interface (i.e., each physical connection to the switch whether between the node and user terminals under its care or between the node and other network nodes) and 2) the aggregate number of calls in the establishment phase. Two additional metrics are the number of buffers used by Setup messages and processor utilization.

The first of the primary metrics, that of the per-interface Setup message arrival rate, can be measured from the layer 2 function. The layer 2 SSCOP mechanism looks at the message header and reads the message type field. If the message type indicates a Setup message, then an arrival message counter can be incremented. The per-interface setup arrival rate is measured by counting the number of setup messages received on an interface over a given interval. At the beginning of each new interval, the arrival message counter is cleared.

The layer 3 call processing mechanism can be monitored to determine the aggregate number of calls in the establishment phase. The call processing module increments an aggregate call establishment counter when a Setup message is processed therein. The call processing module determines when a call has completed the establishment phase by the receiving the appropriate signaling message and at that point

-11-

decrements the aggregate call establishment counter. Though the counter mechanisms are not specifically illustrated in **Figure 3**, one of ordinary skill in the art could readily incorporate such functionality either in hardware or software or use pre-existing counting mechanisms to implement the methodology discussed herein. These are the basic rules for determining setup congestion. More detailed rules for these metrics are discussed below. Also discussed below is the basic control instituted when the metrics fall above or below various thresholds which the discarding of any Setup messages.

Figure 4 illustrates a methodology for tracking the aggregate number of calls in the establishment. Referred to as calls "in progress" these are calls that are being established but have yet to connect and transfer user data. The aggregate number of calls in progress is one of the two major metrics used in instituting call setup congestion control. As shown in **Figure 3**, a mechanism for aggregating the number of calls in progress would derive its data from the call processing function of the layer 3 elements within the SVC controller. The call processing function interprets the exact content of the message and initiates particular processor activity within the controller to carry out any actions that the call message dictates. As mentioned above, a feature of an embodiment of the invention is to not only establish controls, but establish them based on what net effect

-12-

a particular type of message would have on the system. The aggregate number of calls can be tracked using an aggregate calls counter which increments or decrements based upon the content of the message.

For a given potential call, before a call establishment is initiated, the call is an idle state 430. Certain messages, namely those labeled 405, 415 and 425 trigger a transition of state. The establishment phase 400 begins with the receipt of a Call Setup message at a particular node shown as a state transition 405. The aggregate calls counter within a node is incremented, as shown in **Figure 4** when a node receives a Call Setup message 405. When a call is connected, a Connect message 415 is received at the node indicating that the call is in the active phase 420. As such, the call is no longer "in progress," but rather has been fully established. As such, the aggregate calls counter is decremented when a transition from establishment phase 400 to active phase 420 is indicated, usually by a Connect message 415.

A call may also no longer be "in progress" if certain release initiation event message(s) such as Release message 425 are received by the node which may lead to the idle state 430 prior to being fully connected. Thus, whenever Release messages 425 are received during the establishment phase then an aggregate call counter is decremented since a transition back to the idle state 430 is occurring.

-13-

For example a destination busy signal would end the establishment phase and send a Release message to the originator, thus triggering the idle state.

The messages received which trigger the incrementing or decrementing of the counter do not need to belong to a particular call since only an aggregate number is being measured. The aggregate calls counter should always remain running so that node is aware of the totality of calls in progress. This measure is useful since it specifies the total volume of calls being established at the node. When combined with a per-interface measure, controls can be instituted to discard new Setup messages when the Setup message arrival rate and aggregate number of calls in progress exceeds certain thresholds.

Figure 5 is a control and message flow diagram of call congestion monitoring and management.

An algorithm 500 for setup traffic congestion management consists of a global and per-interface portions. Algorithm 500 receives at least two input measurements--aggregate number of calls in progress and a per-interface Call Setup message arrival rate. The per-interface setup arrival rate is derived from an SSCOP mechanism 510 while the aggregate number of calls in progress is derived from a call processing mechanism 530. These mechanisms are described above. By comparing these two inputs to a set of thresholds, upper and lower, for each input, certain control

-14-

decisions may be made, such as when to discard or allow Call Setup messages.

Algorithm 500 makes a per-interface decision as to whether there onset or absence of congestion that requires that Setup messages be discarded and feeds this decision back to block 540. If the decision for a particular interface is that there is congestion, the Setup messages are discarded (block 550) and will not be forwarded to call processing mechanism 530. By discarding messages prior to their being forwarded to the call processing mechanism 530, no processing resources are used by those messages. If the decision is that Setup messages are not be discarded (i.e., allowed), then the Setup messages are forwarded to call processing module 530.

The control action to discard Setup messages should be applied between SSCOP 510 and call processing 530 for two reasons: 1) SSCOP sequencing of messages is not adversely affected and 2) the maximum amount of processor resources are freed (since the control is applied at the earliest possible moment prior to call processing).

Another metric that may be employed in algorithm 500 is a measure of processor utilization which can be compared with thresholds to determine whether or not congestion exists. In nodes that do not have an inter-process buffering mechanism, every effort should be employed to intercept and discard Call Setup

-15-

messages prior to being forwarded to call processing module 530.

Figure 6 is a flowchart of call congestion management.

The Call Setup arrival rate can be defined as the number of call setup messages that arrive over a given fixed interval. Accordingly, at the beginning of the interval, a Call Setup counter should be reset to zero (block 610). The setup message arrival rate is then monitored (block 615) during that interval by incrementing the counter whenever a Call Setup message is encountered. The setup arrival rate measurement is concluded at the end of the interval (block 620). At the end of interval (block 620), a control decision is reached based upon the setup arrival rate and the ongoing continuous measure of aggregate calls in progress.

For the setup arrival rate, two thresholds, an upper setup threshold and lower setup threshold, are set by the system and may be predetermined or dynamically modified depending on desired network behavior or quality of service. Likewise, an upper calls in progress threshold and a lower calls in progress threshold is set by the system for the purposes of comparison with the measure of the aggregate number of calls in progress.

-16-

To make the control decision, a comparison of the setup arrival rate with an upper setup threshold is performed (block 625). If the setup arrival rate is greater than the upper setup threshold and the aggregate number of calls in progress exceeds the upper calls in progress threshold (compared at block 650), then this indicates the onset of congestion (block 655) that must be remedied. At this point, a control is initiated at the SSCOP or buffer level that discards pending (???) and any setup messages that may arrive during the next interval. After the control is enforced, the counter for setup arrival rate is reset and the next interval begun (block 610). At the end of the next interval, if both the setup arrival rate is lower than the lower setup threshold (block 630) and the aggregate calls in progress falls below the lower calls in progress threshold, then the congestion will be deemed abated (block 640).

In the foregoing specification, the invention has been described with reference to specific exemplary embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

-17-

CLAIMS:

What is claimed is:

1. A method of managing congestion in a network controller connecting to a plurality of interfaces comprising:
 - monitoring the utilization and potential utilization of resources in a due to message traffic received by said network controller; and
 - instituting controls to limit further message traffic where said measurements exceed upper thresholds.
2. A method according to claim 1 comprising:
 - releasing said controls where said measurements fall below lower thresholds.
3. A method according to claim 2 wherein said monitoring includes:
 - monitoring the rate at which setup messages arrive at a given interface supported by said network controller; and
 - monitoring the number of aggregate calls in progress over all interfaces supported by said network controller.
4. A method according to claim 3 wherein monitoring the rate at which setup messages arrive includes:

-18-

initializing a setup message arrival counter at the beginning of a interval; and

counting the number of setup messages received at said each interface during said interval.

5. A method according to claim 4 wherein counting includes:

incrementing A setup message arrival counter whenever a setup message is received at said each interface during said interval.

6. A method according to claim 3 wherein monitoring the number of aggregate calls in progress includes:

incrementing an aggregate calls-in-progress counter when a transition for a call from the idle state to the establishment phase is indicated; and

decrementing said aggregate calls-in-progress counter when a transition for a call from the establishment phase to active phase is indicated.

7. A method according to claim 6 wherein monitoring the number of aggregate calls in progress further includes:

decrementing said aggregate calls-in-progress counter when a transition for a call from the establishment phase to idle phase is indicated.

8. A method according to claim 3 wherein the instituting of controls comprises:

-19-

comparing at each interface said setup message arrival rate with an upper setup threshold;

if said setup message arrival rate exceeds said upper setup threshold, then comparing said number of aggregate calls in progress with an upper calls in progress threshold; and

if said number of aggregate calls in progress exceeds said upper calls in progress threshold, then discarding any further setup messages arriving at said each interface before said messages are processed where said setup message arrival rate exceeds said upper setup threshold.

9. A method according to claim 8 wherein the releasing of controls comprises:

if said setup arrival rate does not exceed said upper setup threshold, then comparing said setup arrival rate with a lower setup threshold;

if said setup arrival rate falls below said lower setup threshold, then comparing said number of aggregate calls in progress with an lower calls in progress threshold; and

if said number of aggregate calls in progress falls below said lower calls in progress threshold, then allowing any further setup messages arriving at said each interface to be processed by said controller where said setup message arrival rate exceeds said upper setup threshold.

-20-

10. A method according to claim 9 wherein the instituting of controls includes:

allowing setup messages to be processed by said controller at each interface in any other instance.

11. A method according to claim 3 wherein said aggregate number of calls in progress is monitored from a call processing layer implemented in said network controller and said per-interface setup message arrival rate is monitored from a Service Specific Connection Oriented Protocol (SSCOP) layer implemented in said network controller.

12. A method according to claim 11 wherein said controls are instituted upon setup messages before messages are forwarded from said SSCOP layer to said call processing layer.

13. A method according to claim 3 wherein said interfaces define physical connections between the node in which said controller resides and other nodes connected said node in which said controller resides and define physical connections between the node in which said controller resides and user terminals belonging to said node in which said controller resides.

14. A method according to claim 1 wherein said network controller operates to manage switched virtual circuits.

-21-

15. A method according to claim 3 wherein said setup messages are Call Setup type messages for the establishment of a call in a connection-oriented network.

16. An apparatus including a network switching controller capable of supporting a plurality of interfaces, said apparatus comprising:

a monitoring mechanism adapted to measure the utilization of resources at a given interface and through said network controller; and

a control mechanism coupled to said monitoring mechanism, said control mechanism adapted to control the congestion through said network controller based on signaling by said monitoring mechanism.

17. An apparatus according to claim 16 wherein said network switching controller includes:

a service-specific connection-oriented protocol (SSCOP) mechanism coupled to provide a first signal to said monitoring mechanism, said first signal indicating that a setup message is being received at one of said interfaces; and

a call processing mechanism coupled to said SSCOP mechanism to receive messages forwarded by said SSCOP mechanism, said call processing mechanism further coupled to said control mechanism to provide a second signal indicating that a call is being established through said network controller

-22-

18. An apparatus according to claim 17 wherein said control mechanism is adapted to discard setup messages prior to forwarding to said call processing mechanism for a given interface upon the monitoring mechanism indicating the onset of congestion.

19. An apparatus according to claim 18 wherein said control mechanism is adapted to allow setup messages prior to forwarding to said call processing mechanism for a given interface upon the monitoring mechanism indicating the absence of congestion.

20. An apparatus according to claim 18 wherein said onset of congestion is indicated by a count of the number of first signals received in a given interval if that count exceeds an upper setup message threshold and if a count of the number of second signals received over all interfaces exceeds an upper calls-in-progress threshold.

21. An apparatus according to claim 19 wherein said absence of congestion is indicated by a count of the number of first signals received in a given interval if that count falls below a lower setup message threshold and if a count of the number of second signals received over all interfaces falls below a lower calls-in-progress threshold.

-23-

22. An article comprising a computer readable medium having instructions which when executed manages congestion in a network controller having a plurality of interfaces, said instructions causing:

monitoring the utilization and potential utilization of resources in a due to message traffic received by said network controller; and

instituting controls to limit further message traffic where said measurements exceed upper thresholds.

23. An apparatus for managing congestion in a network controller having a plurality of interfaces comprising:

means for monitoring the utilization and potential utilization of resources in a due to message traffic received by said network controller; and

means for instituting controls to limit further message traffic where said measurements exceed upper thresholds.

24. An apparatus according to claim 23 comprising:

means for releasing said controls where said measurements fall below lower thresholds.

25. An apparatus according to claim 24 wherein said means for monitoring includes:

-24-

means for monitoring the rate at which setup messages arrive at a given interface supported by said network controller; and

means for monitoring the number of aggregate calls in progress over all interfaces supported by said network controller.

26. An apparatus according to claim 25 wherein means for monitoring the rate at which setup messages arrive includes:

means for initializing a setup message arrival counter at the beginning of a interval; and

means for counting the number of setup messages received at said each interface during said interval.

27. An apparatus according to claim 26 wherein means for counting includes:

means for incrementing a setup message arrival counter whenever a setup message is received at said each interface during said interval.

28. An apparatus according to claim 25 wherein means for monitoring the number of aggregate calls in progress includes:

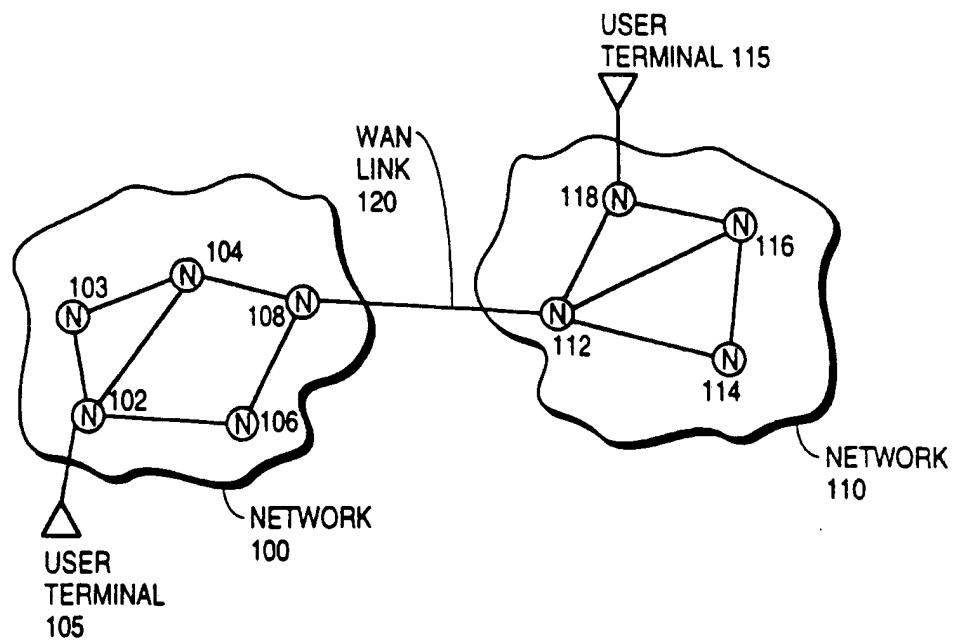
means for incrementing an aggregate calls-in-progress counter when a transition for a call from the idle state to the establishment phase is indicated; and

-25-

means for decrementing said aggregate calls-in-progress counter when a transition for a call from the establishment phase to active phase is indicated.

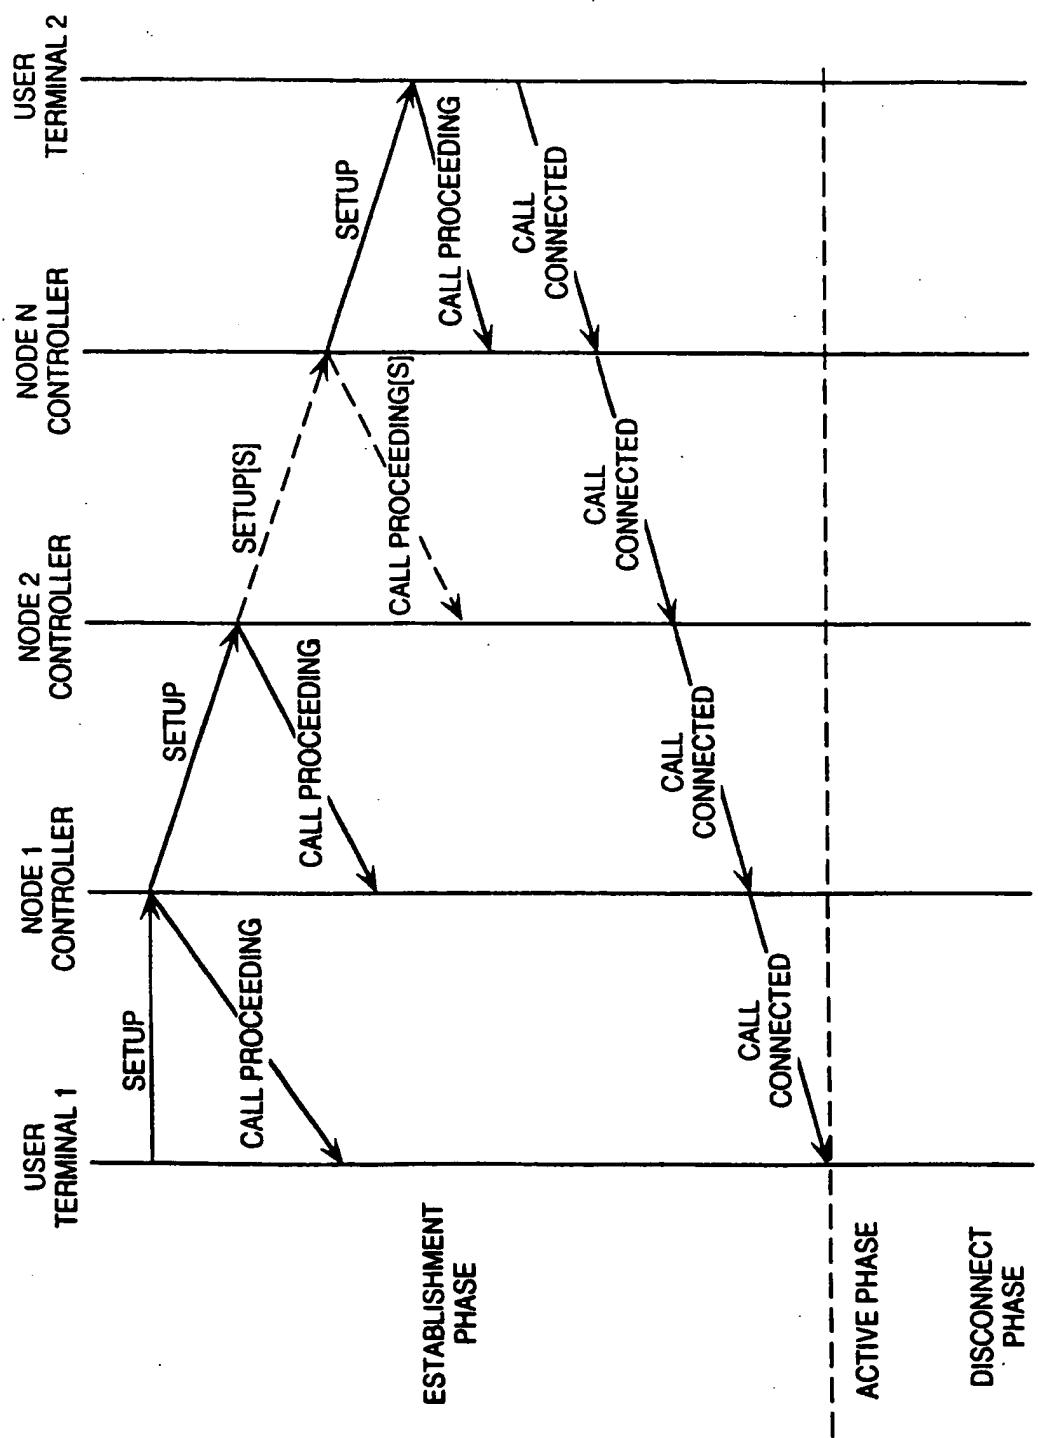
29. An apparatus according to claim 28 wherein means for monitoring the number of aggregate calls in progress further includes:

means for decrementing said aggregate calls-in-progress counter when a transition for a call from the establishment phase to idle phase is indicated.



N=NODE

Fig. 1

**Fig. 2**

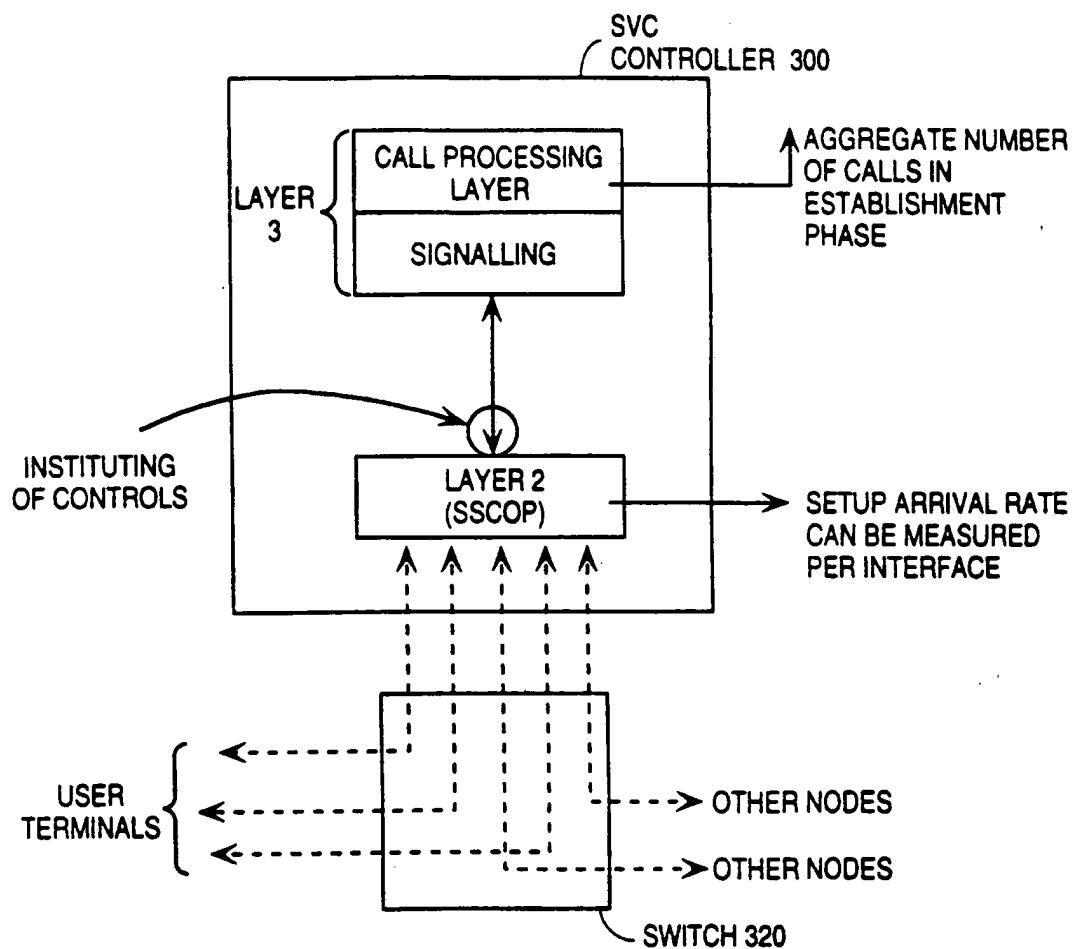


Fig. 3

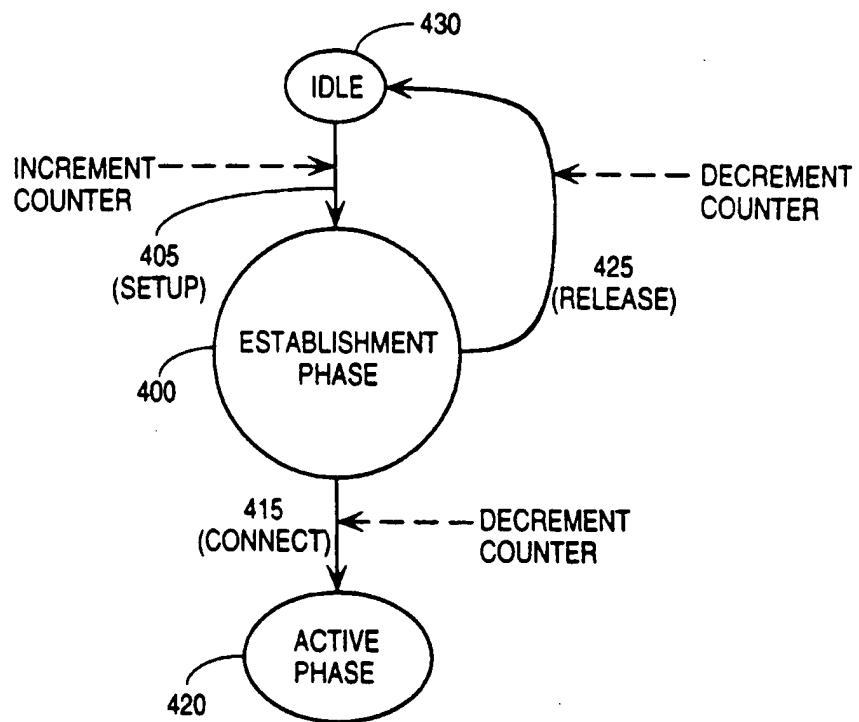


Fig. 4

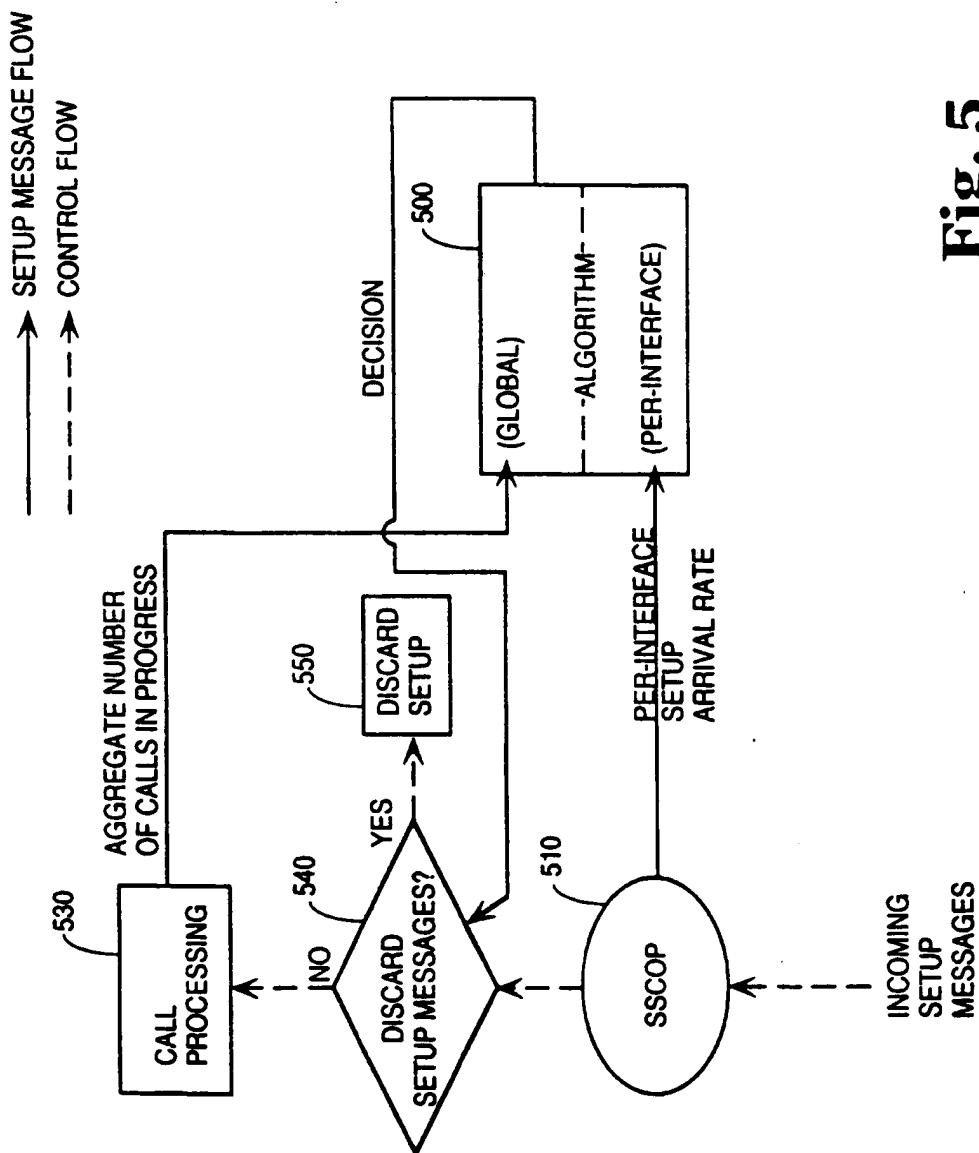


Fig. 5

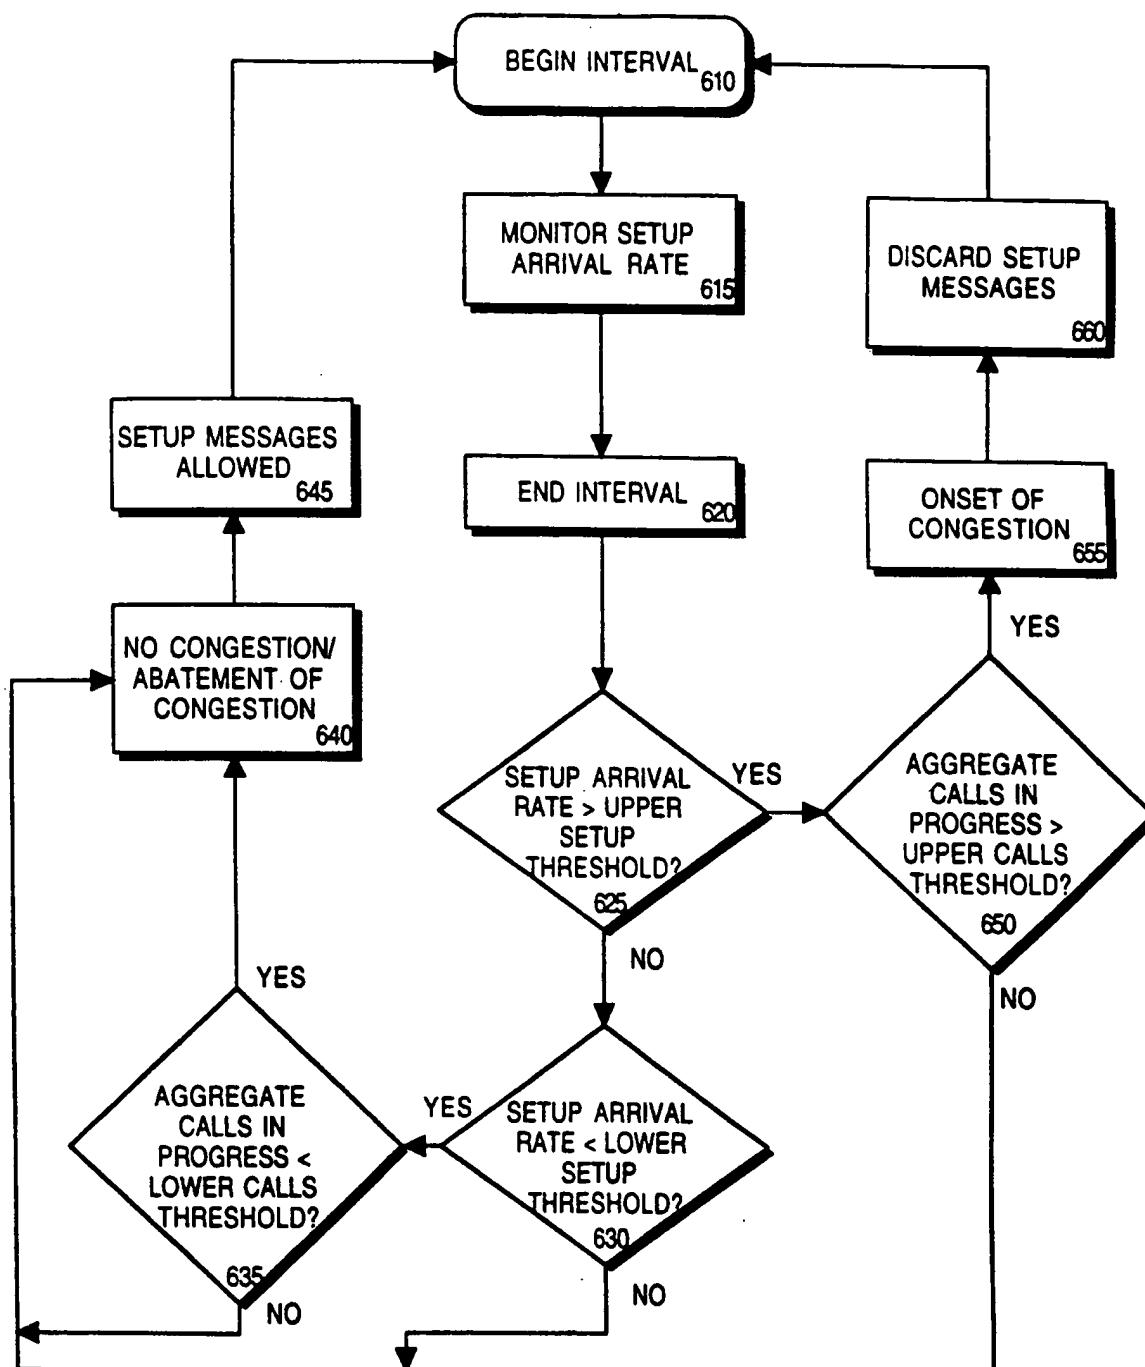


Fig. 6

INTERNATIONAL SEARCH REPORT

Int'l. Application No
PCT/US 00/02604

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04Q11/04 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04Q H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	MITO M ET AL: "B-ISDN signalling protocol processing for large multiplexed subscriber system" 1995 IEEE INTERNATIONAL CONFERENCE ON COMMUNICATIONS. CONVERGING TECHNOLOGIES FOR TOMORROW'S APPLICATIONS. ICC '96. CONFERENCE RECORD (CAT. NO.96CH35916), PROCEEDINGS OF ICC/SUPERCOMM '96 - INTERNATIONAL CONFERENCE ON COMMUNICATIONS, DALLAS, TX, USA,, pages 663-668 vol.2, XP002140624 1996, New York, NY, USA, IEEE, USA ISBN: 0-7803-3250-4 figures 2,8 page 663, column 2, paragraph 2 page 664, column 1, line 31 -column 2, line 8 page 666, column 1, paragraph 4.2 -page 667, column 1, line 6	1,14, 16-18, 22,23
Y		2,24
A		3-5, 25-27
		-/-

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

20 June 2000

Date of mailing of the international search report

04/07/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl.
Fax: (+31-70) 340-3016

Authorized officer

Lamadie, S

INTERNATIONAL SEARCH REPORT

Inte.	onal Application No
PCT/US 00/02604	

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	page 667, column 1, paragraph 4.3 US 5 473 604 A (SHAW ROBERT F ET AL) 5 December 1995 (1995-12-05) figures 2-6	1,14,16, 22,23 3,4,8, 11,13, 15,17, 18,25-27
A	column 2, line 38 - line 47 column 3, line 3 - line 32 column 5, line 4 - line 11 column 5, line 24 - line 27 column 6, line 16 - line 39 column 7, line 5 - line 22 column 7, line 59 - line 61 column 8, line 44 - line 53 -----	
Y	US 5 519 690 A (SUZUKA TETSUYA ET AL) 21 May 1996 (1996-05-21) figure 4 abstract -----	2,24

INTERNATIONAL SEARCH REPORT

Information on patent family members

Inte. .onal Application No

PCT/US 00/02604

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5473604	A 05-12-1995	NONE	
US 5519690	A 21-05-1996	JP 7177196 A	14-07-1995